REMARKS

Applicants respectfully request reconsideration of this application. Claims 1, 2, 4, 5, 7-14, 17, 19-32, and 70-83 remain in this application. Claims 1, 14, 70, and 76 have been amended. No claims have been added.

Examiner Interview

Applicant wishes to thank the Examiner for the courtesy of the telephonic interview on June 15, 2009, in which Applicant's counsel and the Examiner discussed the rejection of the claims and potential amendments. No agreement of patentability was reached.

Rejections under 35 U.S.C. § 112, first paragraph

Claims 1-2, 4-5, 17, 19-32, and 70-83 stand rejected under 35 U.S.C. § 112, first paragraph as failing to comply with the written description requirement. In particular, the Examiner asserts that the limitation "paths that a series of three or more of said optical network devices connected by links" is not adequately described in the specification. Applicant respectfully submits that this limitation is supported Figures 2 and 4A-4D and paragraphs 75-77 and 81.

Furthermore, the Examiner asserts that the limitation "wherein said databases are not specific to or created responsive to path request messages" is not adequately described in the specification. "Any negative limitation or exclusionary proviso must have basis in the original disclosure ... Note that a lack of literal basis in the specification for a negative limitation may not be sufficient to establish a *prima facie* case for lack of descriptive support." (MPEP, § 2173.01(i)). Applicant believes this limitation is definite as a basis for this claim limitation is found in Figures 15-19 and paragraphs 157-183, in which Applicant's network topology databases are created using the distributed search technique and not in response to a request to a lightpath. Furthermore, Applicant discloses in Figure 20 and paragraphs 185-188 disclose that the system allocates a

lightpath in response to receiving message that indicates a request for a path/wavelength combination.

Rejections under 35 U.S.C. § 103(a)

Claims 1, 2, 4, 7, 11-12, 14, 17, 20, 24-27, 70, 73-74, and 76-79 stand rejected under 35 U.S.C. § 103(a) as being obvious in view of Ho et al. ("A Novel Distributed Control Protocol in Dynamic Wavelength-Routed Optical Networks", IEEE Communications Magazine, November 2002), Smith et al., U.S. Patent No. 7,171,124, Roorda, U.S. Patent Publication No. 2002/0186432, and Desnoyers, U.S. Patent Applicant No. 2002/0186432.

Ho discloses dynamically allocating an optical path from a source node to a destination node (Ho, Abstract, p.38). The source node is provided with a routing table that defines all possible paths to the different destinations nodes (Ho, p.39, 2nd column). The paths of the routing table are known to the source node because the paths are defined offline (Ho, p. 39, 2nd column). Ho's path allocation scheme dynamically updates the link status of multiple paths from the source node to the one destination node (Ho, p. 39, 2nd column). The link status of each link of a node on the paths is determined by either periodically sending probe messages along the paths or sending the probes messages in response to a path allocation request (Ho, p. 39, 2nd column). Assigning the proper path involves selecting the path and selecting the proper wavelength along the path (Ho, p. 38, 1st column). Because all of the routing between the source and destination nodes is predefined in the routing table, path selection just involves wavelength selection (Ho, p. 39, 2nd column). Wavelength selection selects the best lightpath between the source and destination nodes (Ho, p. 39, 2nd column). This is done by determining critical links (i.e. paths with high traffic), and broadcasting to other nodes to avoid the critical links during path selection (Ho, p. 40, 2nd column). Nevertheless, Ho does not disclose the organization of the routing database, other that the database contains paths between source and destination nodes.

Smith discloses a system that routes and switches information on a dense wavelength division multiplexing network (DWDM) path between a source and destination node (Smith, Abstract). They system routes and switches information

through multiple nodes in the network (Smith, Fig. 1, Col. 4, lines 13-22). Regenerating nodes ("regenerators") provide optical-electric-optical (OEO) wavelength conversion and/or regeneration in the network core (Smith, Col. 4, lines 34-36). The system further includes a network topology database to store topology data of the network (Smith, Col. 6, lines 65-66). In response to receiving a path request message, requesting a path from one source node to one destination node, the system finds a number of paths available between these specific source and destination nodes (Smith, Col. 6, lines 12-16). The system builds a search tree of these source-destination paths from the topology data stored in the database (Smith, Col. 6, lines 65-66). The search tree is organized by regenerator levels from the source node to other nodes along the paths to the destination node (Smith, Figure 5A, Col. 19, line 58 – Col. 10, line 11). Using this search tree, the system further selects viable paths in the search tree groups the paths by the number of regenerators on each path and stores these selected paths (Smith, Col. 6, lines 26-33). In particular, regenerator placement module groups the paths into 'm' sets by the number of 'k' regenerators that are on the path (Smith, Fig. 4A, step 63, Fig. 5B, Col. 10, lines 53-67). In the example illustrated by Fig. 5B, there is one set with one path having one regenerator (k=1), a second set of three paths with two regenerators (k=2), and a third set of one path having one regenerator (k=3) (Smith, Fig. 5B, Col. 10, lines 54-59). In addition, Smith discloses computing the costs of paths (Smith, Fig. 6B, Col. 13, lines 49-54). If a path is selected, the routing module "destroys the stored paths," as these paths are no longer needed (Smith, Figure 3B, step 20, Col. 8, lines 57-60). However, Smith does not disclose the organizational structure of the network topology database nor does Smith suggest using the paths stored in memory as a replacement for the topology database. Thus, Smith discloses three different structures that could be interpreted as a network database: (1) a network topology database that Smith fails to disclose the information stored in or the organizational structure of the network topology database; (2) the search tree of paths from the source node to different destination nodes; and (3) paths stored in memory from one source node to one destination node, wherein these paths are stored in response to receiving a path request message.

Roorda discloses a photonic network that provides separation of passthrough channels that form drop channels at the input of a switching node (Roorda, Abstract). A

node of the network has the ability to switch an input fiber to an output fiber as well as adding/dropping traffic with the photonic network (Roorda, Figure 3A, paragraph 52). In addition, Roorda discloses discovering topology information by each node recognizing its own connections to neighboring nodes (Roorda, paragraph 106). Roorda does not disclose how its topology information is stored or updated in a database.

Desnoyers discloses a computer that discovers a network topology and uses this topology to make routing decisions in a network comprising computers and switching nodes of an electrically switched network (Desnoyers, Abstract). Desnoyers' computer comprises only one network interface (Desnoyers, Fig. 2, col. 13, lines 11-13). This computer processes data and transfers data to other computers (Desnoyers, Col. 3, lines 46-48). However, each computer is an end point in the network and does not switch data from other computers (See, e.g., Desnoyers, Fig. 1). Instead, Desnoyers discloses switching nodes that switch data from computers and other switching nodes (Desnoyers, Col. 3, lines 58-61).

Desnoyers' computer, but not the switching node, determines the network topology in an iterative fashion by sending out identification request messages specifically to different other computers and switches in the network (Desnoyers, Col. 5, lines 10-20). Each newly discovered switch responds with an identification response message identifying the port the identification request message was received on and the path the identification request message traveled to that switch (Desnoyers, Col. 6, lines 48-53). The requesting computer uses the received responses to generate network topology information (Desnoyers, Col. 5, lines 3-9).

In addition, Desnoyers discloses that the computer builds the network topology information because the switching nodes do not have the requisite processing and storage resources to determine and store a network topology (Desnoyers, Col. 2, lines 39-50). The computer uses this information to configure one or more of the switching nodes (Desnoyers, Col. 10, lines 35-37). Because Desnoyers' computer is gathering topology information and configuring one or more of the switching nodes, Desnoyers' computer is acting as a network management system for Desnoyers' switching nodes.

Desnoyers' response message includes a path traveled from the computer to the switching node that sent the response message (Desnoyers, Col. 7, lines 11-20). The path

stored in these messages includes the input/output ports and node identification for each intermediate switching node as well as the identification and input port for the switching node that sent the response message back to the originating computer (*Id.*). For example, switching node 11(6) (SN6) sends a response message back to computer 12(1). The sent response message contains the path information "<P1, SN1, P2 | P3, SN6>." P1, P2, and SN1 are the input port, output port, and switching node identification of the intermediate switching node 11(1). P3 and SN6 are the input port and switching node identification of the sending switching node 11(6) that sent the response message (*Id.*). Thus, the switching node SN6 sends the response message that includes information of the path that the identification request message traveled to SN6 and no further.

Furthermore, Desnoyers does not disclose the structure of its topology database.

Applicant respectfully submits that the combination of Ho, Smith, and Roorda do not teach or suggest Applicant's claims. (1) Applicant respectfully submits that the combination does not teach or suggest creating and maintaining a network topology database responsive to receipt of messages that include end to end path identification information transmitted from an optical network device in that end to end path that is coupled to a destination optical network device of that end to end path. (3) The combination does not teach or suggest organizing databases such that paths are grouped by reachable destination node with at least two different groups of available paths to two different reachable destination nodes.

Creating and Maintaining

Applicant respectively submits that none of Ho, Smith, Roorda, and/or Desnoyers teach or suggest creating and maintaining a network topology database responsive to receipt of messages that include end to end path identification, wherein each of said messages include end to end path information transmitted from a node in that end to end path that is coupled to the destination node of that end to end path and an end to end path comprises a series of three or more nodes. Ho discloses receiving link status messages that merely identify the status of a link between two nodes. Because Ho's link status only comprises a link between two nodes and not an end to end path that comprises a series or

three or more nodes, Ho cannot teach or suggest maintaining a network topology database by receipt of messages that include end to end path identification, wherein each of said messages include end to end path information transmitted from a node in that end to end path that is coupled to the destination node of that end to end path and an end to end path comprises a series of three or more nodes.

Smith does not disclose updating the set of paths stored in memory that is created in response to a path request. Quite the opposite of maintaining, Smith discloses that this set of paths is destroyed as soon as the path request is fulfilled. Furthermore, Smith discloses that a signal and control system that updates Smith network topology database. However, Smith does not disclose how that database is updated by the signal and control system. In addition, Smith fails to disclose maintaining the search tree. Therefore, Smith does not teach or suggest maintaining a network topology database by receipt of messages that include end to end path identification, wherein each of said messages include end to end path information transmitted from a node in that end to end path that is coupled to the destination node of that end to end path and an end to end path comprises a series of three or more nodes.

Furthermore, because Roorda does not disclose how to update the topology information stored in a database, Roorda does not teach or suggest maintaining a network topology database by receipt of response messages that include end to end path identification, wherein each of said messages include end to end path information transmitted from a node in that end to end path that is coupled to the destination node of that end to end path and an end to end path comprises a series of three or more nodes.

Desnoyers discloses the destination node sending the path information, but not from a node that is coupled to the destination node. Because Desnoyers discloses the destination node sending the path information, and not a node coupled to the destination node, Desnoyers does not teach or suggest maintaining a network topology database by receipt of response messages that include end to end path identification, wherein each of said messages include end to end path information transmitted from a node in that end to end path that is coupled to the destination node of that end to end path and an end to end path comprises a series of three or more nodes.

For example, claim 1, as amended, requires "receipt of messages that each carry end to end path identification information back to the originating access nodes transmitted from an optical network device in that end to end path that is coupled to a destination optical network device of that end to end path, wherein the end to end paths include paths that are a series of three or more of said optical network devices connected by links on which a set of wavelengths is available for establishing a lightpath ...said optical network devices acting as access nodes each creates and maintains, responsive to receipt of the messages that include the end to end path identification information, a network topology database, which represents the network topology for that access node, through sustained storage of the possible end to end paths, with costs, from that access node to all other reachable destination nodes."

Furthermore, claim 14, as amended, requires "each of said paths is represented in said database by a cost and a series of three or more nodes and interconnecting links over which that path travels ... a database module that creates and maintains said database responsive to receipt of messages that each include end to end path identification information transmitted from a node in that end to end path that is coupled to a destination node of that end to end path, wherein said database is not specific to or created responsive to path request messages."

In addition, claim 70, as amended, requires "an available path comprising a sequence of three or more nodes ... creating and maintaining said database responsive to receipt of messages that each include end to end path identification information transmitted from a node in that end to end path that is coupled to a destination node of that end to end path, and wherein said databases are not specific to or created responsive to path request messages."

Claim 76, as amended, requires "an end to end path is a series of three or more of nodes connected by links on which a set of wavelengths is available for establishing a lightpath ... each of the received responsive connectivity messages include end to end identification information transmitted from a node in that end to end path that is coupled to a destination node of that end to end path."

The above quoted limitations are not described or suggested by Ho, Smith, Roorda and/or Desnoyers. While there are various uses for the invention as claimed,

several such uses are discussed at Figures 5 and 18 and paragraphs 68, 69, 88-91, and 174. Thus, while the invention is not limited to the uses discussed on these pages, it should be understood that Ho, Smith, Roorda, and Desnoyers do not enable these uses and the above quoted limitations do.

Grouping Paths Based on Common Destination Nodes

The Examiner admits that Ho does not teach or suggest a "grouping paths based on common destination nodes" and relies on Smith to supply the missing element (Office Action, p. 3). In support of his rejection, the Examiner states that Smith "teaches in FIG. 5B and 6B grouping paths with common destination nodes for evaluating and selecting the best path for routing" (emphasis added, *Id.*). However, Applicant respectfully submits that this set of paths disclosed in Figs. 5B and 6B are for a single destination node and Smith does not teach or suggest a grouping based common destination nodes.

Furthermore, the Examiner admits that Smith fails to disclose grouping of paths to common destination nodes, but states that it is obvious "that the same method can be applied a plurality of times to find multiple paths from one source node to a plurality of destinations" (*Id.*). However, this is not what Applicant claims in claim 1, 14, 70, and 76. These claims are directed towards network topology databases that are organized such that possible paths to each of to reachable destinations are grouped together under that reachable destination node with at least two different groups of available paths to two different reachable destination nodes. Because the Examiner misstated what Applicant is claiming, independent claims 1, 14, 70, and 76 are not rendered obvious by the combination.

In addition, because neither Roorda nor Desnoyers disclose the structure of the topology databases, neither reference can teach or suggest the claimed element.

Moreover, it would not be obvious to one of skill in the art to modify Ho, Smith, Roorda, and/or Desnoyers to have possible end to end paths to reachable destinations grouped together under that reachable destination node with at least two different groups of available paths to two different reachable destination nodes. In order to support an obvious rejection, the Examiner must show that the difference between the prior art cited and the claimed invention would have been obvious to one of skill in the art (Fed. Reg.

Vol. 72, No. 195, p. 57528). The Examiner asserts that it would obvious to apply the method of Ho, Smith, Roorda, and Desnoyers "a plurality of times to find multiple paths from one source node to a plurality of destinations" (Office Action, p. 4). As above, the Examiner apparently equates Smith's single group of source-destinations paths to Applicant's database having <u>possible end to end paths to reachable destinations grouped together under that reachable destination node</u> with <u>at least two different groups of available paths to two different reachable destination nodes</u>. Applying Smith's path selection method a plurality of times at one access node would result in a plurality of transient independent source-destination sets of paths stored in memory. Each set of paths is created solely in response to separate requests for a path. The sets of paths are then destroyed after that request is fulfilled. However, these sets of transient, independent sets of paths stored in memory would not be considered a database to one of ordinary skill in the art. A database is "a collection of organized, related data" (Random House Webster's' College Dictionary, 2nd Ed., 1997). While each individual set of paths maybe organized within that given set, multiple instances of these sets of paths are not organized relative to each other. Each set of paths is independent and disconnected from the other sets of paths. These sets of paths are only created is response to a request for one source-destination path. Thus, Applicant respectfully submits that applying Smith path selection method a plurality of time does not teach or suggest Applicant's database having <u>possible end to end paths to reachable destinations grouped together under that reachable destination node</u> with <u>at least two different groups of available paths to two different reachable destination nodes</u>

Furthermore, it would not be obvious to one of skill in the art to modify Smith to organizing multiple paths to multiple destinations with the multiple paths grouped by common destination, because Smith actually teaches away from this limitation. This is because Smith actually discloses destroying the set of paths to one destination stored for path allocation once one of those paths has been allocated (Smith, Figure 3B, step 20, Col. 8, lines 57-60). Thus, as Smith actually teaches away organizing multiple paths to multiple destinations, with the multiple paths grouped by common destination, it would not be obvious to one of skill in the art to modify Smith.

In addition, organizing a topology database with possible end to end paths to reachable destination under that destination would not have been obvious modification to one of skill in the art. As above, the Examiner must show that the difference would have been obvious to one of skill in the art. One example of one of skill in the art can be found in the well-known Open Shortest Path First (OSPF) protocol. In OSPF, the topology database is organized as a link state database for every node known in the network (See, e.g., Moy, "OSPF Version 2", RFC-2328, pp. 18-21; and Applicant's specification, paragraph 16). With this database, a node generates a shortest path first (SPF) tree representing paths from the source node to other destinations (Moy, p. 21-23). When used for optical networking, this SPF tree is at the individual link and lambda levels because there are multiple lambdas per link and different lambdas may provide different characteristics (Applicant's Specification, paragraph 16). Using this SPF tree, the node selects a path to route (Moy, p. 21). It should be noted that the since OSPF is a protocol typically used in optically networking to create topology databases, Moy would be considered one of skill in the art. Furthermore, even though it was known at the time of Moy that a database may be organized in a different fashion, Moy (and other who use OSPF) chose to organize the topology database as a link/lambda state database and/or a link/lambda SPF tree, instead of a topology database with paths grouped together under each destination. Thus, because Moy organized the OSPF database(s) differently than Applicant's topology database and Moy is one of skill in the art, it would not have been obvious to one of skill in the art to try and organize the topology database with paths grouped together under each destination.

Furthermore, organizing a topology database with paths group to a destination leads to results that would not have predicted. A demonstration of a modification that leads to predictable result can be basis for obviousness (Fed. Reg., Vol. 72, No. 195, 57529, Rationales (A)-(D)). However, a topology database organized by paths grouped to a destination allows a more efficient search for paths than a database and/or SPF tree as used with OSPF. OSPF uses $O(N^2)$ algorithms to search for path, whereas a path grouped topology can be search for a path using $O(\log N)$ algorithms, with N being smaller in the latter case for the same size network (Applicant's Specification, paragraphs 128-129). For example, the number of operations to search an SPF tree representing a network of 10

nodes, eight fibers/node, and 40 wavelengths/fiber is approximately 3210^2 (*Id.*). In contrast, searching the same network using the path grouped topology database requires approximately log(640) operations, which is much more efficient than for OSPF (*Id.*). Thus, this topology database organization is much more efficient than others known in the art. Therefore, a topology database with paths group to a destination leads to results that would not have predicted.

Thus, none of Ho, Smith, Roorda, or Desnoyers teach or suggest <u>possible end to end paths to reachable destinations that are grouped together under that reachable destination node and that there are at least two different groups of available paths to two different reachable destination nodes.</u> Nor is it obvious to one of skill in the art to modify Ho and/or Smith to have such a topology database.

For example, claim 1, as amended, requires "<u>optical network devices</u> acting as access nodes each creates and maintains, responsive to receipt of the messages that include the end to end path identification information, a network topology database, which represents the network topology for that access node, through sustained storage of the possible end to end paths, with costs, from that access node to all other reachable destination nodes, each of said paths in said databases having associated with it the wavelengths available on that path, <u>wherein the possible end to end paths in each of said databases are organized such that all available paths to each of the reachable destination nodes are grouped together under that reachable destination node, wherein at least one of said databases stores at least two different groups of available paths to at least two different reachable destination nodes</u>."

Claim 14, as amended, requires "a destinations structure to store each of said destinations in a single entry, <u>each of said destination's single entry to reference paths to that destination</u>, wherein <u>each of said paths is represented in said database by a cost and a series of three or more nodes and interconnecting links over which that path travels, and wherein said database stores at least two different groups of paths to at least two different reachable access nodes</u>..."

Claim 70, as amended, requires "...locating a reachable destination node in a structure of a database, wherein <u>said structure stores a non-duplicative set of the plurality of destination nodes in the optical network, wherein each of said plurality of reachable</u>

destination nodes in the structure references a group of available paths to that destination node, an available path comprising a sequence of three or more nodes and interconnecting links of those of the available paths that lead to that destination node and the all of the available paths in each group sorted at least in part by cost, each such available path having associated to it the set of one or more available wavelengths along that path to here; and wherein said database stores at least two groups of the available paths to at least two different reachable destinations…"

Claim 76 requires "…responsive to receiving the response connectivity messages, storing in a database the collected end to end paths, the collected end to end paths organized in the database such that all available paths to each of the destination nodes are grouped together under that destination node, and the database stores at least two groups of the available paths to at least two different destinations, wherein said database is not specific to or created responsive to path request messages…"

The above quoted limitations are not described or suggested by Ho, Smith, Roorda and/or Desnoyers. While there are various uses for the invention as claimed, several such uses are discussed at Figures 3-4 and paragraphs 85-90. Thus, while the invention is not limited to the uses discussed on these pages, it should be understood that Ho, Smith, Roorda, and Desnoyers do not enable these uses and the above quoted limitations do.

With regards to Applicant's claim 14, Applicant respectfully submits that the combination does not teach or suggest a separately stored "destinations structure to store each of said destinations in a single entry" and that "each of said destination's single entry to reference paths to that destination." The Examiner does not specifically state where these limitations are disclosed in Ho, Smith, and/or Roorda. Furthermore, none of Ho, Smith, Roorda, and/or Desnoyers teach or suggest specifically detail a destination structure as claimed in claim 14. For example, because Ho, Roorda, and Desnoyers fail to disclose the structure of Ho's routing table Roorda's/Desnoyers' topology database, respectively, none of Ho, Roorda, and Desnoyers teach or suggest this limitation. As per above, because (1) Smith's structure of Smith's topology database is not disclosed, (2) Smith's search is not organized by common destination, and (3) Smith's stored set of

paths are for only one destination, Smith cannot teach or suggest the claimed limitation. Furthermore, because this type of database structure enables a more efficient search for paths (see above), using a destinations structure as claimed is not an obvious modification to the combination. Thus, the combination does not render claims 14 and claims 17 and 19-32 that depend on it.

For at least these reasons, Applicant respectfully submits that the claims discussed above are allowable. The Applicant respectfully submits that the additional dependant claims are allowable for at least the reason that they are dependent on an allowable independent claim.

Claims 14, 17, 20, and 24-27 stand rejected under 35 U.S.C. § 103(a) as being obvious in view of Ho, Smith, Roorda, and Chaudhuri (U.S. Patent No. 7,039,009).

Chaudhuri discloses a method for lightpath restoration in an optical network (Chaudhuri, Abstract). Each node in the optical network determines the network topology using a routing protocol, such as Open Shortest Path First (OSPF), that propagates link state advertisements throughout the entire network (Chaudhuri, Col. 7, lines 3-9). Using OSPF, each node in the network eventually determines the entire network topology and not a topology specific to that node (Chaudhuri, Col. 7, lines 9-10). Chaudhuri further discloses that the network topology comprises an address for each node at each end of each link; a total number of active channels on each link; a number of allocated channels on each link; a number of preemptable channels on each link; a number of reserved restoration channel on each link; Shared Risk Link Groups throughout the reconfigurable optical network; and optional physical layer parameters for each link (Chaudhuri, Col. 25, lines 53-62).

Applicant respectively submits that <u>none of Ho, Smith, Roorda, and/or Chaudhuri teach or suggest creating and maintaining a network topology database responsive to receipt of messages that include end to end path identification, wherein each of said messages include end to end path information transmitted from a node in that end to end path that is coupled to the destination node of that end to end path and an end to end path comprises a series of three or more nodes</u>. As per above, none of Ho, Smith and Roorda

teach or suggest this limitation. Because Chaudhuri relies on link state advertisement to determine network topology, Chaudhuri does not teach or suggest this claim limitation.

For at least these reasons, Applicant respectfully submits that the claim discussed above is allowable. The Applicant respectfully submits that the additional dependant claims are allowable for at least the reason that they are dependent on an allowable independent claim.

Claims 8-10, 21-23, 71-72, and 80-83 stand rejected under 35 U.S.C. § 103(a) as being obvious in view of Ho, Smith, Roorda, Desnoyers, and Ho et al., ("A Framework for Service-Guaranteed Shared Protection in WDM Mesh Networks", IEEE Communications Magazine, February 2002) ("Ho2"). Claims 5 and 19 stand rejected under 35 U.S.C. § 103(a) as being obvious in view of Ho, Smith, Roorda, Desnoyers, and Deo et al., "Graph Theory with Applications to Engineering and Computer Science"). Claim 13 and 28-29 stand rejected under 35 U.S.C. § 103(a) as being obvious in view of Ho, Smith, Roorda, Desnoyers, and Jukan et al. ("Constraint-Based Path Selection Methods for On-Demand Provisioning in WDM Networks", IEEE INFOCOM, 2002). Claims 30 and 32 stand rejected under 35 U.S.C. § 103(a) as being obvious in view of Ho, Smith, Roorda, Chaudhuri, and Moy ("OSPF Version 2", RFC 2328, IETF, April 1998). Claim 31 stand rejected under 35 U.S.C. § 103(a) as being obvious in view of Ho, Smith, Roorda, Desnoyers, and Pulkkinen et al., U.S. Patent Publication No. 2003/0172356. Claim 15 stands rejected under 35 U.S.C. § 103(a) as being obvious in view of Ho, Deo, Smith, Roorda, Desnoyers, and Date ("An Introduction to Database Systems" by C.J. Date, Addison-Wesley 1986). Claim 75 stands rejected under 35 U.S.C. § 103(a) as being obvious in view of Ho, Ho2, Smith, Roorda, Desnoyers, and Shami ("Performance Evaluation of Two GMPLS-Based Distributed Control and Management Protocols for Dynamic Lightpath Provisioning in Future IP Networks", IEEE, 2002).

All of the above claims depend from one of the above identified independent claims. It is respectfully submitted that the above identified cited references, individually or in combination, fail to disclose or suggest the limitations set forth the above independent claims.

Invitation for a telephone interview

The Examiner is invited to call the undersigned at 408-720-8300 if there remains any issue with allowance of this case.

Charge our Deposit Account

Please charge any shortage to our Deposit Account No. 02-2666.

Respectfully submitted,

BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN

Date: June 15, 2009

Eric Replogle
Reg. No. 52, 161

1279 Oakmead Parkway
Sunnyvale, CA 94085-4040
(408) 720-8300